

# SCIENCE INFORMATION SYSTEMS NEWSLETTER

[Plain Text Site](#) | [Planet Animation \(4.5Mb\)](#) | [Copyright & Credit](#) | [Nasa Privacy Statement](#)

[CURRENT ISSUE](#) | [ARCHIVES](#) | [LINKS](#) | [SISN SEARCH](#)

## *Office of Space Science*

### Science Information Systems Program

Read the latest on NASA's information systems

[www-sisn.jpl.nasa.gov](http://www-sisn.jpl.nasa.gov)

53

# Using A Distributed Inventory System to Archive Science Data

**J. Steven Hughes and Susan K. McMahon, Planetary Data System, Jet Propulsion Laboratory**

The advent of the World Wide Web (Web) and the ability to easily put data repositories on-line has resulted in a proliferation of digital libraries. The heterogeneity of the underlying systems, the autonomy of the individual sites, and distributed nature of the technology has made both interoperability across the sites and the search for resources within a site major research topics. The Planetary Data System (PDS) addresses both issues using standard Web protocols and meta-data labels to implement an inventory of on-line resources across a group of sites. The success of this system is strongly dependent on the existence of and adherence to a standards architecture that guides the management of meta-data within participating sites.

PDS [1] is an active science data archive managed by scientists for NASA's planetary science community and has been in operation since 1990. Envisioned as a long-term archive, the PDS early on emphasized the development of a standards architecture that would include both the science data and the meta-data necessary for interpreting diverse data storage formats as well as understanding the context under which the data were captured. This standards architecture has been used to create a high quality, peer-reviewed, science data archive of about five terabytes that is stored on compact disk (CD) media. The meta-data in this archive, even though collected to ensure the usability of the science data for future scientists, has also allowed the majority of the archive to be made available on-line via the Web as a digital library. As a component of this digital library, the Distributed Inventory System (DIS) was developed to identify all resources across the heterogeneous, autonomous, and distributed nodes of the PDS. This inventory includes all archived data sets, pending data sets, and any resources that support the use of data sets. The DIS is a lightweight solution to managing access to PDS resources, taking advantage of the wealth of meta-data available in the archive.

## Developing a science data archive

In 1983, the Committee on Data Management and Computing (CODMAC) issued a report that set the guidelines for the development of a science data archive. [2] This report resulted from the observation that a wealth of science data would ultimately cease to be useful and probably be lost if a process was not developed to ensure that the science data were properly archived. In particular, the report proposed two major goals. First, the committee recommended that the data be transferred to more stable media. Second, the committee

recommended that sufficient ancillary and meta-data be captured and archived with the data to ensure that future users would be able to understand how to interpret the data formats, as well as understand the context under which the data was collected and processed.

Additionally, the report suggested that sophisticated searches for data sets were important. In particular the scientists wanted the capability to find data sets through relationships with other entities. For example, using the relationships between data sets, spacecraft, and instruments, scientists wanted the ability to identify images that had been captured using specific filters on a specific camera type on any spacecraft.

The PDS was organized in 1985 and five years of design and development followed. It went on-line in 1990 with about 20 data sets in its archive. A major component of the system was a high-level data set catalog that allowed the searching and ordering of any data set in the archive through an on-line interface of about 88 user views. To ingest data into the archive, the PDS developed a data ingestion procedure that includes a formal peer review. This review involves a peer review committee comprised of peer scientists who review the science data for validity and usability and technical staff that review the collected science data and meta-data for adherence to the standards architecture. The PDS today has about 400 peer-reviewed data sets in its archive. Another 100 are pending peer review. The standards architecture and procedures, including the peer review, are documented in three volumes, the Data Preparation Workbook [3], the Standards Reference [4], and the Planetary Science Data Dictionary (PSDD) [5]. The PDS data model and the meta-data development and management process have been described in [6] and [7].

Within the planetary science community an example of a science data set is the collection of about 50,000 Mars images returned by the Viking Orbiter spacecraft in 1976. An individual image within this data is an example of a data set granule or product.

## Data management approach

The original vision for managing PDS data focused on adhering to a standards architecture, distributing the data for management by discipline scientists, and allowing access to the data through an integrated set of catalogs.

The PDS is a distributed system, consisting of five science discipline nodes, two support nodes and a central node. Contracted out in most cases to universities, the discipline nodes manage the science data within their discipline while also providing science expertise. Originally, each node was solely responsible for distributing their data. The central node managed the data set catalog and forwarded orders to the discipline nodes. The data were distributed on CD or on tape.

The standards architecture is the cohesive element of the PDS, allowing the autonomous and distributed science discipline nodes to produce archive volumes that are very much alike in organization, quality, and usability.

With the advent of the Web and a growing volume of data to ingest into the archive, the PDS had to evolve. In particular, Web technology has allowed advanced on-line interfaces to be developed in a fraction of the time that was previously required. This allowed better access to the archive also but resulted in more users. Also, additional copies of the data could be put on-line at other sites with relative ease. This addressed bandwidth problems but required more management because of redundancy. Finally, users wanted access to data sets that were not fully archived. Because only archived data sets were ingested into the data set catalog, pending data sets were accessible only to a select few that were told where to find the data.

This combination of new technology and the need to manage multiple on-line copies of archived and pending data sets and non-data resources, led to the collection of requirements for an inventory system. Limited resources dictated a lightweight solution focusing on simple and automatic operations.

## Purposes of DIS

There were originally three purposes for the DIS. First, it was to improve user access to the PDS by identifying what data sets and resources were available either on or off-line at the time of the query. To accomplish this, an inventory was to be taken across all nodes on a regular basis to track the location of all copies of a data set, verify the on-line status on a regular basis, and check Web site status. This system was to build on the existing data set catalog of archived products.

The second purpose was to help the user by clarifying the quality of the available data sets. Because the ingestion process can take months and in some cases years to complete, it was important to notify the user of the quality of the data. To accomplish this, the `archive_status` and `archive_status_note` attributes were created.

The third purpose was to provide disaster recovery. With a complete inventory of all data sets and resources resident in a single database, a backup copy of the database could be made and the inventory could be brought up at an alternate site.

## The DIS approach

The DIS consists of two major components, the central DIS manager and individual DIS component at each of the nodes. All the components are conceptually identical. The

DIS manager contains the inventory across the entire PDS while a node component contains the inventory local to that node. The inventory consists primarily of object labels.

## Object labels

The DIS inventory primarily consists of a DIS label for each object in the inventory. Currently the PDS inventory includes all fully archived data sets, most pending data sets, and non-data resources. Examples of resources include Web sites, user search aids such as catalogs, jukeboxes, utility software, and even people. Additional ancillary labels exist for servers and sites. In fact, new object types can be added as needed by simply designing a new label.

A DIS label describes an object using meta-data in key-word-value format as shown in Figure 1. For a data set object, the keywords and values used are primarily those already specified in the PDS standards for describing a data set in the archive. In particular, a data set label includes identification information (`data_set_id` and `data_set_name`), status information (`archive_status` and `archive_status_note`), on-line links to the data (`link` and `volume_link`), relational information (`instrument_name`, `target_name`), and text descriptions (`description`, `terse_description`.) The `node_name` keyword identifies discipline nodes that distribute the data set and `curating_node_id` identifies the node responsible for creating the DIS label. It is important to note that for an archived data set, the `description` keyword in the DIS label simply has a link to an entry in the data set catalog. When entering the data set catalog, the data set description is the first information displayed.

```
OBJECT = DATA_SET;
LABEL_HISTORY_NOTE = 1998-04-3 DSUPD LN:AY:DN;
DATA_SET_NAME = VO1/VO2 MARS VISUAL IMAGING ...
DATA_SET_ID = VO1/VO2-M-VIS-2EDR-BR-V2.0;
ARCHIVE_STATUS = ARCHIVED;
ARCHIVE_STATUS_NOTE = NONE;
DATA_SET_DESC = http://pds.jpl...
DATA_SET_TERSE_DESC = http://pds.jpl...
DATA_OBJECT_TYPE = IMAGE;
START_TIME = 1976;
STOP_TIME = 1980;
NODE_NAME = IMAGING;
CURATING_NODE_ID = IMAGING;
TARGET_NAME = DEIMOS, MARS, PHOBOS, STAR;
MISSION_NAME = PRE-MAGELLAN, VIKING;
INSTRUMENT_HOST_NAME = VIKING ORBITER 1 ...
INSTRUMENT_NAME = VISUAL IMAGING SUBSYSTEM ...
VOLUME_ID = VO_1001, VO_1002,...
KEYWORDS = IMAGING, MARS, VIKING, VO_1001, VO_1002 ...
LINK = OFFLINE;
VOLUME_LINK = ftp://pdsimage.wr.../cdroms/vo_1001/,
              ftp://pdsimage.wr.../cdroms/vo_1002/, ...
END_OBJECT;
```

Figure 1 - Data Set DIS Label

The resource label identifies non-data resources that are available from one or more nodes. Again as seen in Figure 2, identification information, on-line links, status, and a text description are provided. The status can be used to notify users of any planned changes in status or current conditions that might interest the user.



```

OBJECT = RESOURCE;
LABEL_HISTORY_NOTE = 1998-01-01 GEO, ...
NAME = Mars Navigator;
DESC = The Mars Geoscience Navigator provides the capability
      to locate, display, download, and order geoscience data
      products from Mars missions.;
NODE_NAME = GEOSCIENCES;
CURATING_NODE_ID = GEOSCIENCE;
STATUS = UP;
DATA_SET_NAME = VO1/VO2_MARS_VISUAL_IMAGING_SS ...
KEYWORDS = CAMERA, DEIMOS, IMAGE, IMAGING, MARS, ...
LINK = http://window.wustl.edu/marsnav/;
END_OBJECT;

```

Figure 2 - Resource DIS Label

A site label is created for each on-line site within the system. As seen in Figure 3, this label provides a site identifier, location information, a link, and a description. Also, since there is a status and since the links provided in data set or resource label contains site identifiers, the status of the site can be merged with a link and displayed as a result of a user query.

```

OBJECT=SITE;
NAME=pdssbn.astro...;
DESC=PDS Small Bodies Node Web Server;
SERVER_HOST_ID=129.2...;
SERVER_HOST_NAME=pdssbn.astro...;
STATUS=UP;
LINK=http://pdssbn.astro.../;
END_OBJECT;

```

Figure 3 - Site DIS Label

## Object label creation

The labels for archived data sets as shown in Figure 1 are created automatically from information contained in the data set catalog database. In particular, the relational information such as instrument\_name and target\_name are produced by simple joins. The keyvalues keyword however, is a collection of values from any keywords in the label that could be useful for searching.

The labels for pending data sets and resources are created manually by the curating node. However, for resource labels, the values for the key values keyword are augmented automatically by the system. Using the identifiers provided for the data\_set\_id keyword, a utility extracts all related information from the data set catalog and augments the values for the key values keyword.

## Object class definition

The object classes are defined using a special DIS label. Figure 4 illustrates the class definition for the resource object. Each keyword has a text definition, a data type, sequence number for results ordering, and additional formatting information.

```

CLASS = RESOURCE;
OBJECT = The object element identifies the objects class (type).
CHAR 10 F STD;
NAME = The name element provides an identifier for the RESOURCE.
CHAR 20 FSM STD;
DESC = The desc element provides the description of the RESOURCE.
TEXT 30 FSM FREE;
LINK = The link element provides the HTTP protocol string for ...
TEXT 41 FSM LINK:PROMPT Resource Online;
KEYVALUES = The keyvalues element identifies target, missions, ...
CHAR 70 KFS STD;
LABEL_REVISION_NOTE = The label_revision_note element provides ...
TEXT 92 F STD;
LABEL_LINK = The label_link element provides the HTTP protocol ...
TEXT 93 FSM LINK:PROMPT Source Label;
...
END_CLASS;

```

Figure 4 - Resource Object Class Definition Label

## DIS system architecture

The current implementation of the DIS is based on the Meta-Data List Server (MEDALIS) engine, a common gateway interface (CGI) script written in PERL. For the PDS implementation the engine has been named pdsserv. As seen in Figure 5 the pdsserv module accesses three databases, object.db which contains the DIS labels, class.db which contains the class description labels, and system.db which contains system information. As a cgi-script, the module accepts a query in the form of an HTTP protocol compliant message. This allows users to query the module directly through the location field in their favorite browser.




Figure 5 - DIS manager system architecture

The components of the central DIS manager are a pdsserv module, class.db, system.db, and an object.db that contains a complete inventory of the PDS, as is illustrated in Figure 5. At each of the discipline nodes, the components are identical except that the object.db database contains only the labels for objects local to that node. To update the DIS manager database, a polling script is periodically run to query each node for its inventory and the results are concatenated and processed to produce the object.db for the DIS manager.

## DIS server engine

The pdsserv module, (MEDALIS engine), accepts queries in the form of HTTP protocol compliant queries. The two primary requirements for this query engine are the ability to constrain the query using any keyword in any label in the object.db database and that the value of any keyword can be returned. Figure 6 provides a synopsis of a MEDALIS query. As can be seen, constraints are specified using any number of keyword=value clauses. The values, which can include regular expression components, are compared with the values of the specified keywords for every label in the object.db database to produce the result set.

### SYNOPSIS

```
http://pds/cgi-bin/pdsserv.pl?  
  KEYWORD=VALUE*           * => zero or more  
  &RETURN=KEYWORD*         * => zero or one.  
  &FORMAT=FORMAT**
```

### CLAUSES

KEYWORD=VALUE - search constraint  
RETURN=KEYWORD - KEYWORD selection (default->all)  
RETURN=HELP - returns OBJECT class description  
RETURN=DATA - redirects to LINK if query result is one  
FORMAT=FORMAT - format specification  
& - AND logical connective  
, !, ( ) - OR, NOT, and nesting available when encoded

### FORMATS

FORMAT=HTML -> HTML list format - Default  
FORMAT=HTML:M -> HTML list format, minimum keywords  
FORMAT=TABLE -> HTML table format  
FORMAT=PLAIN -> ODL text format  
FORMAT=DISDB -> DIS database dump format

Figure 6 - Query Synopsis

Similarly, the keyword values to be returned are specified using any number of return=keyword clauses. If no return clauses are given, all keywords in the selected labels are returned. An optional format clause allows the results to be returned in an HTML list format, HTML table format, semicolon delimited table format, or DIS label format. The logical connectives, AND, OR, NOT and parenthetic nesting are allowed. Entering HELP as a query returns the query synopsis.

Being object-based, the MEDALIS engine allows the return of the three modes of an object, the class description, object description, and the object itself. Since the DIS label is itself the object description, a standard query returns full or partial descriptions of one or more objects. The class description of an object is returned by including the clause RETURN=HELP. Note that the returned objects would have to be of the same type for this query to be valid. Finally, the actual object being described by the label can be returned by appending RETURN=DATA to the query. This clause causes a redirection to the location specified in the LINK=location statement in the label. This displays the information at that link or what the label designer considers to be the actual object.

Figure 7 shows several MEDALIS queries. The first query specifies that only data set labels that have "VG" as the first

characters of the data\_set\_id are to be returned. For the PDS inventory, all labels describing VOYAGER data sets would be returned. The second query would return all JUPITER related data sets by searching the keyvalues keyword. Note that the first query represents a database attribute search capability where the values of a single keyword are being searched. The second query searches the values of the key-values keyword - a concatenation of values from several keywords - and is more flexible. However the search is still constrained to a controlled set of values.

- 1) pdsserv.pl?OBJECT=DATA\_SET&DATA\_SET\_ID=VG  
returns Voyager data sets
- 2) pdsserv.pl?OBJECT=DATA\_SET&KEYVALUES=JUPITER  
returns Jupiter data sets
- 3) pdsserv.pl?OBJECT=RESOURCE&FORMAT=TABLE  
returns resources in table format
- 4) pdsserv.pl?OBJECT=DATA\_SET&RETURN=HELP  
returns data sets class description
- 5) pdsserv.pl?OBJECT=DATA\_SET&DATA\_SET\_ID=DIM&RETURN=DATA  
redirects to LINK

Figure 7 - Example Queries

The third query simply returns a list of all RESOURCES in the inventory in HTML table format. The fourth query, by specifying RETURN=HELP returns the class description for data set objects. The final query searches for a data set label that contains "DIM" in its data\_set\_id but redirects the server to return the information at the location specified in the LINK keyword of the DIM data set label.

In Figure 8, one of the six results from the first query in Figure 7 is displayed in HTML list format. Not all keywords in the data set label are shown in the result since the class description for the data set object has specified a limited set of keywords to be returned as the default. DATA\_SET\_NAME provides identification information and the ARCHIVE\_STATUS specifies that the data set fully archived. Three link keywords in particular should be noted. The LINK keyword provides the location for what the label creator considers the entire data set. The data set in this case does not exist as a single entity that can be pointed to - such as a Web page - so is considered offline.

The VOLUME\_LINK keyword however provides the location of 24 on-line volumes. The site status has been determined by accessing the appropriate SITE label and merging in the value of the STATUS keyword. Finally, the LABEL\_LINK is a link that can be used to force the engine to return the entire label.

```
DATA_SET  
DATA_SET_NAME = VG1/VG2 JUPITER IMAGING SCIENCE SUBSYSTEM EDITED EDR ...  
DATA_SET_TERSE_DESC = This data set contains compressed level-2 (unprocessed) images from  
the Voyagers 1 and 2 encounters with Jupiter. It also contains documentation, software, and index  
directories to support access to the image files in the volume set.  
LINK = OFFLINE  
ARCHIVE_STATUS = ARCHIVED  
NODE_NAME = IMAGING  
VOLUME_LINK = Volume Online : vg_0006 - site: www.pdsimage.jpl.nasa.gov - site_status: UP  
Volume Online : vg_0007 - site: www.pdsimage.jpl.nasa.gov - site_status: UP ...  
LABEL_LINK = Source Label - site_status: UP
```

Fig. 8 - Query Results

Figure 8 - Query Results

## DIS interfaces

Developers can create HTML forms as client interfaces to the pdsserv module. A very simple interface is illustrated in Figure 9. This form allows the user to constrain queries for either data sets or resources, choose a format for the results, and accepts a list of search keywords. The query produced by this form searches the keyvalues keyword in each label.

Figure 9 - HTML Interface

## PDS access hierarchy

As mentioned previously, the majority of the PDS archive exists on CD-ROM media. Proceeding with a library metaphor, this represents the corpus of a traditional library. However, in addition to the actual documents, the meta-data necessary for identifying, searching, describing, and using the documents has also been packaged with the documents and placed in the library.

By placing the archive resident on CD media into a juke-box and making it accessible from the Internet, the PDS becomes a digital library accessible by the science community, the educational community, and the general public. In 10, this primitive level of access is represented as the base of the Figure. In some cases, the CD content is moved to disk for improved efficiency.

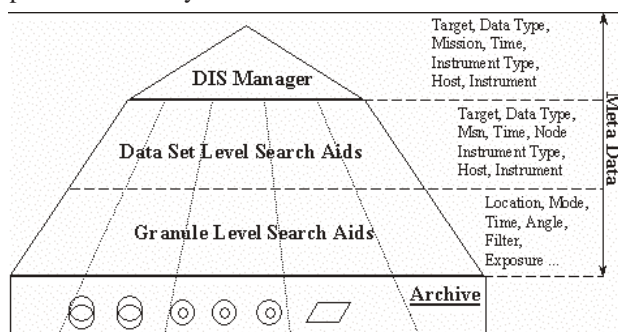


Figure 10 - PDS Access Hierarchy

At the next level - for popular data sets and when resources are available - granule level search aids are developed. For example, the Planetary Image Atlas, a search aid developed by the PDS Imaging node, allows for users to search for

images within selected image data sets. This is accomplished by extracting the meta-data stored with each image in the Viking Orbiter Mars Image data set and loading it into a relational database system. Web interfaces are then customized to allow the search for any image by using map and forms based interfaces. In fact, for the Viking data set, a user could compose a complex query for any of the approximately 30 attributes associated with an image, such as camera filter, center point location, or exposure duration.

At the data set level, the data set catalog allows for the search of archived data sets using more global attributes such as data set start\_time, stop\_time, data\_set\_id, and data\_set\_name. Relationships to spacecraft, instruments, missions, and targets can also be used to identify data sets.

As Figure 10 shows, the DIS manager tops off the hierarchy. It is dependent on the data set catalog for the inventory of archived data sets and the discipline nodes for the inventories of pending data sets and resources.

## Current status

The DIS has met the three original requirements for the system. The current DIS system is operational with labels for all archived data sets, pending data sets, and available resources. Through the simple HTML form interface, the system is in constant use by scientists, the general public, discipline node personnel, and PDS management. We have several discipline node components up and are implementing the automatic update process. It has been observed that the frequency of update at some of the nodes does not warrant automatic upload, so some updates are simply e-mailed to the central node for directly updating the DIS manager database. This flexibility is allowed. Depending on available funding, more sophisticated user interfaces will be built and modifications will be made to the pdsserv engine to make it more efficient.

Other uses have also been found for the DIS. For example, in PDS operations it has been useful for reconciling the data set catalog and individual node inventories. Management is also considering producing status reports for the data set ingestion queue using archive\_status values.

A simple data set granule or product server could also be implemented by creating labels from meta-data stored with the products. This would allow a user or program to access a product by simply specifying the data\_set\_id and product\_id. Since the product type is included in the meta-data, existing utilities could be integrated to allow the transformation and display of the products.

The DIS is also being considered as a component for building interoperability between PDS and other science data systems within the NASA community. The DIS, using standard Web protocols, can be queried from external systems to locate on-line data and resources within the PDS. The DIS is also being considered as a lightweight inventory solution for other distributed data systems.



## Future direction

Future directions focus on staying compatible with advances in digital library technology. In particular, the developers are considering the use of XML as the language for the DIS label. The availability of XML editors will preclude the need for developing smart editors for the current DIS labels. In addition, DIS query results in XML format would be more portable, an important requirement for system interoperation. Query results in XML format can also handle more complex results and can be displayed in more sophisticated formats.

The MEDALIS engine has been converted to JAVA and a graphical interface has been developed. This is currently being tested and will soon be made operational.

The PDS DIS provides a first level of interoperability across a heterogeneous, autonomous, and distributed data system. The DIS is a lightweight solution that was developed using standards Web protocols, a simple query language, and labels that describe objects using simple keyword-value statements. It provides the user with the ability to search and display any object in the inventory. Query results provide users with descriptions of the objects and links on-line locations. On-line status information is also displayed.

This system, even though significant on its own as a simple yet powerful tool for developing interoperability across distributed sites, is also an important case study in the development and management of meta-data for a digital library. It is primarily because of the availability and consistency of the meta-data in the archive that this system was a success. The populating of the database simply consisted of extracting meta-data from the archive and reformatting it into the DIS label format. The success of this system is strongly tied to the standards architecture developed for and adhered to by the PDS.

*This paper was originally presented, along with a poster exhibit, at the Advances in Digital Libraries conference this past spring. See the Report on Advances in Digital Libraries in Issue 52*

*<<http://www-sisn.jpl.nasa.gov/ISSUE52/adl.html>>. The research described here was carried out by the Jet Propulsion Laboratory, California Institute of Technology, under a contract with the National Aeronautics and Space Administration.*

Learn more about the Planetary Data System and its Distributed Inventory System by emailing Steve Hughes <[Steve.Hughes@jpl.nasa.gov](mailto:Steve.Hughes@jpl.nasa.gov)> or accessing the PDS web site <[http://pds.jpl.nasa.gov/pds\\_home.html](http://pds.jpl.nasa.gov/pds_home.html)>.


## Acknowledgments and References

The authors wish to acknowledge Elizabeth Duxbury (PDS Image Node) and Todd King (PDS PPI Node) as prime contributors in the initial discussion that led to the development of the DIS, as well as the PDS staff in general and Mike Martin in particular for continued development of the PDS standards architecture

1. Russell, C. T., et al., Special Issue: "Planetary Data System, Planetary and Space Science," Pergamon, Vol. 44, No. 1, 1996.
2. Arvidson, R.A., et al., "Issues and Recommendations Associated with Distributed Computation and Data Management Systems for the Space Sciences," National Academy Press, 1986.
3. Planetary Data System Data Preparation Workbook, JPL Internal Document, JPL D-7669, Part 1, Jet Propulsion Laboratory, 1993
4. Planetary Data System Standards Reference, JPL Internal Document, JPL D-7669, Part 2, Jet Propulsion Laboratory, 1995
5. Planetary Science Data Dictionary, JPL Internal Document, JPL D-7116, Jet Propulsion Laboratory.
6. Hughes, J.S., Li, Y.P., "The Planetary Data System Data Model," Proceedings of Twelfth IEEE Symposium on Mass Storage Systems, 1993, 183-189.
7. Hughes, J. S., McMahon, S. K., "The Planetary Data System - A Case Study in the Development and Management of Meta-Data for a Scientific Digital Library," Lecture Notes in Computer Science, Second European Conference on Research and Advanced Technology for Digital Libraries, 1998, 335-350. ■

# *Welcome to Antarctica! A Tour of Images from RADARSAT*

*Judy Laue, Raytheon ITSS, Goddard Space Flight Center*



Before and after images showing the disintegration of the Antarctic Peninsula ice shelves. Scientists captured the first image using the ERS-1 satellite in 1992 and the second image using RADARSAT in 1997, revealing major changes to the coastline. In 1978, scientists predicted the disintegration of these ice shelves due to global warming.

**G**oddard Space Flight Center's (GSFC) Scientific Visualization Studio (SVS) has produced "Welcome to Antarctica!" a World Wide Web-based virtual tour of images of Antarctica revealing features of this remote continent that

have never before been seen. The NASA-launched Canadian satellite RADARSAT created high resolution radar maps of the farthest reaches of the South Pole that the SVS visualizers used to create the news-breaking visualizations.

Scientists from Ohio State University's Byrd Polar Research Center and animators from the SVS designed the virtual tour of the southernmost part of the Antarctic. The tour begins and ends at McMurdo Station, and includes images and movies of valleys, hills, ice shelves, ice streams, glaciers, and snow dunes. The web site also includes images of the first flight over the South Pole as well as movies of the 1911 race to the Pole and the perilous 1899 voyage of the Belgica. In addition, the web site includes descriptions of the images and movies.

The SVS assists scientists with visualizing data using traditional and leading-edge tools and techniques.

Learn more about the Antarctica and the SVS by Emailing the SVS Director Dr. Horace Mitchell at <horace.mitchell@gsfc.nasa.gov>. Also, take the virtual tour of the Antarctica  
<<http://svs.gsfc.nasa.gov/imagewall/antarctica.html>>. ■



# Implementing MM5 on Goddard's Computing Systems: a Performance Study

**John Dorband and Udaya Ranawake, Goddard Space Flight Center, Jules Kouatchou, Morgan State University - Baltimore, Maryland, and John Mikalakes, Argonne National Laboratory**

One of the objectives of the Earth Space Data Computing Division (ESDCD) as part of the Earth Sciences Enterprise at Goddard Space Flight Center (GSFC), is to make the computational capabilities of scalable vector and parallel computers available to scientists involved in Earth Sciences research. ESDCD supports initiatives to report benchmark results of various climate and weather forecast models in order to indicate how well its NASA Center for Computational Science (NCCS) computers perform in comparison with other machines. In addition, ESDCD aims to provide these results to the scientific community and to create a user-friendly environment so that users of NCCS systems can easily access, run, and examine performance of climate and weather simulation codes. To this end, the team proposes to make the fifth-generation PSU/NCAR Mesoscale Model MM5 [1, 2] available to users on NCCS facilities.

This article presents the results of a performance study of MM5 on the Cray J90, the Cray T3E and a cluster of PCs (theHIVE) at GSFC. Using a small and a large size test problem, the efficiency of MM5 on these computers is analyzed. In particular, the elapsed times, the speedup, the number of floating point operations per second are reported. In addition, for each platform, the ratio of cost to performance is determined.

## Description of MM5 code

MM5 is limited-area weather model designed to simulate mesoscale and regional-scale atmospheric circulation over regions ranging from several thousand to several hundred kilometers. It is a primitive-equations model employing finite differencing for the computation of atmospheric dynamics (advection and diffusion) under both hydrostatic and nonhydrostatic assumptions, both of which employ time splitting. The model includes a multiple-nest capability, a four-dimensional data assimilation capability, a variety of physics options. MM5 is maintained in the public domain by NCAR and is used both internally and by more than three-hundred groups at institutions worldwide. It is used for real-time forecasting, climate, storm research, studies related to air-quality and urban heat islands, and simulation to support basic atmospheric research.

The official version of MM5 was designed and maintained for efficiency on vector and shared-memory parallel architectures [3], as well as on workstations (SGI, SUN, IBM, DEC, HP, and PCs). Prior to 1998, however, the model did not run on message-passing distributed memory (DM) parallel computers. Computational demands for very fine resolution and large problem domains motivated the development of a distributed memory parallel option [5] to MM5. MM5 is parallelized using two-dimensional data domain decomposition. Two- and three-dimensional data arrays containing state variables (wind velocity, temperature, moisture, and pressure) as well as diagnostic and intermediate fields are partitioned in two dimensions and the resulting subdomains are distributed over processors. Extra memory is allocated at processor subdomain boundaries for ghost points to store off-processor operands for finite differencing and horizontal interpolation [4]. This is implemented using the Runtime System Library (RSL) [7] and the Fortran Loop and Index Converter (FLIC) [6]. RSL provides domain decomposition, local address space computation, distributed I/O, and inter-processor communication supporting parallelization of both the solver and mesh refinement code. FLIC translates at compile-time to generate a parallelized code (that only the compiler sees) from a single version of the source model. The approach is essentially directiveless, requiring only a small amount of information to direct the translation.

The parallel version was close enough to the original non DM-parallel version to be incorporated and maintained as part of the official code. The DM-parallel option first appeared in Release 8 of MM5, in March 1998, has continued through Release 9, in May 1998. MM5 currently runs on the IBM SP2, Cray T3E, SGI Origin2000, HP Exemplar, distributed memory configurations of DEC Alpha processors, the Fujitsu VPP, and Beowulf-like networks of PCs. A performance analysis on some of these computers appears in [5]. The code is in use operationally on an IBM SP2 at the U.S. Air Force Weather Agency producing twice-daily real-time forecasts over 36 domains world-wide. The DM-parallel MM5 is also in use at the US EPA and a number of other facilities in the US and internationally.

## Target machines

NCCS is the high-performance scientific computing facility operated, maintained, and managed by ESDCD. The NCCS mission is to provide computing resources and support services to enable space and Earth scientists who are currently funded by NASA Headquarters to accomplish their research goals. Among the computers available at NCCS, a Cray J90, the Cray T3E and a cluster of PCs can be mentioned. These machines are targets of the analysis of the MM5 code.

### Cray J90

The Cray J90 system in the study is a shared memory vector computer that has 32 CPUs and each one of them can perform 200 Mflops. It has a 10 nanosecond clock, 8 gigabytes of main memory, an aggregate disk transfer rate of 320 MB/sec, a vector length of 64 and 390 gigabytes of disk capacity and runs the Cray Research, Inc. Unicos operating system.

### Cray T3E

The T3E is a massively parallel processor system which consists of 32 to 2048 Processors Elements (PE). Each PE has a DRAM memory of 64 megabytes to 2 GB. Due to the several programs the T3E contains the user can expect to have only 60 megabytes from each PE. In the T3E, the PE are Alpha microprocessor called DEC chip 21164 capable of 600 million floating point operations per second. Like the Cray J90, the T3E runs the Cray Research Inc. UNICOS operating system. At NCCS the T3E has 1024 processors with 305 GFLOPS and 480 gigabytes disk storage.

### The HIVE

The Highly-parallel Integrated Virtual Environment (theHIVE) is a Beowulf-class parallel computer, i.e. a cluster of PCs running LINUX, connected by its own private LAN. It is a network of 64 compute nodes, BEEs, and 2 types of host or frontend nodes, the QUEEN and one or more DRONES. The QUEEN administers theHIVE and the DRONES allow access to theHIVE by users and data gathering facilities, such as remote sensing satellite antennas. Each node has two 200Mhz Pentium Pro processors (128 processors total), 256K of cache, 14 GB of disk space (900 GB Total). The nodes are interconnected with 100 Mhz fast Ethernet switches.

The purpose of theHIVE is to demonstrate the usefulness of low cost commodity high performance computers on NASA's earth and space science applications and to study the ability of theHIVE to fulfill the need for a high performance computing resource integrated seamlessly into the working environment of multiple personal workstations.

The features of each of these computers are summarized in Table 1. theHIVE is relatively cheaper compared to the other two platforms; the cost/Mflop is about \$8.00 on theHIVE whereas on the Cray J90 and the Cray T3E it is \$95.00 and \$20.00 respectively.

**Table 1: Description of the Platforms**

	Cray J90	Cray T3E	the HIVE
OS	Unicos	Unicos	Linux
Processor	Cray J90se	DEC chip 21164	Pentium Pro
Peak	200 Mflops/proc.	598 Mflops/proc.	200 Mflops/proc.
Clock Speed	100 Mhz	300 Mhz	200 Mhz
Disk	390 GB	1470 GB	900 GB
N	32	1024	128
Cost/proc.	\$19,000.00	\$12,000.00	\$1,600.00

### Performance results

Here, results obtained from two test problems are presented:

#### Test Problem 1

Benchmark case: SESAME  
Domain size: coarse mesh: 25x28x23  
fine mesh: 34x37x23  
Number of time steps: 160 coarse domain  
480 fine mesh  
Size of data sets: 25 Mb

#### Test Problem 2

Benchmark case: Hurricane Opal  
Domain size: 109x119x35  
Number of time steps: 91  
Size of data sets: 43 Mb

Test Problem 1 (Sesame) has two domains (coarse and fine). The coarse and fine domains have 90 km and 30 km resolution, 25 and 34 grid points in the x-direction, 28 and 37 grid points in the y-direction, 24 vertical levels respectively. A 12 hour forecast was done on the Cray J90 and theHIVE.

Test Problem 2 (Hurricane Opal) has one domain with 45 km resolution, 109 and 119 grid points in the x- and y-directions respectively and 36 vertical levels. Here the team executes the MM5 code for a 3 hour forecast (91 time steps) on the Cray J90, the Cray T3E, and theHIVE.

The team is interested in finding (for runs with at most 16 processors) the elapsed time, the speedup, and the number of floating point operations per second when the number of processors varies. These parameters will determine how efficient is MM5 on the target platforms.

The Hurricane Opal case is a problem larger than the Sesame case. On DM computers, the team wants to make sure that the work load is evenly spread and that all the processors are busy most of the time. It is more likely that with a small size problem such as Sesame, such goals will not be achieved.

For the experiments, the shared memory version of MM5 on the Cray J90 will be employed, and the DM-parallel version on the Cray T3E and theHIVE. Both versions are incorporated in the official code and produce the same forecast results.

It is important to mention that the runs on the Cray J90 were done in a dedicated mode, i.e. team members were the only users utilizing the requested processors. On the Cray

T3E, the system runs automatically in a dedicated mode but on theHIVE there is no such mechanism. However, during each of our experiments on theHIVE, the team made sure that members were the only users in the system.

## Sesame test case

Experiments for this test case were done on the Cray J90 and theHIVE only. The elapsed time and the speedup is reported as function of the number of processors. Results are summarized in Figures 1 and 2 respectively. On both systems, it is observed that the elapsed time decreases as the number of processors increases with the best results on the Cray J90. However the speedup curve shows moderate performances on theHIVE. For this problem, the efficiency of theHIVE falls very quickly below 50% (with 8 processors). The Sesame case is not an ideal case for gauging parallel efficiency, though because it's so small. Processors probably run out of work in the first stages of computations.




Figure 1. Sesame: elapsed time (in seconds) as function of the number of processors.




Figure 2. Sesame: speedup as function of the number of processors.

In addition, in Figure 3 the number of floating point operations (Mflops) per second is plotted and in Figure 4 the

cost/Mflop as function of the number of processors. Again, it is noted that MM5 is faster on the Cray J90 compared to theHIVE (Figure 3). However, if the performances of MM5 is examined on both platforms with respect to the cost, theHIVE offers the best result, i.e. the cost/Mflop is lower on theHIVE (Figure 4).




Figure 3. Sesame: number of floating point operations (Mflops) per second as function of the number of processors.




Figure 4. Sesame: Cost/Mflop as function of the number of processors.

## Hurricane opal test case

To continue the analysis, new experiments are performed with a larger problem on the Cray J90, the Cray T3E, and theHIVE. On the Cray T3E, it was not possible to run the algorithm with less than four processors because of the large data set involved.

Elapsed time and speedup as function of the number of processors are plotted in Figures 5 and 6 respectively.






Figure 5. Hurricane Opal: elapsed time (in seconds) as function of the number of processors.




Figure 6. Hurricane Opal: speedup as function of the number of processors.

It is observed that MM5 requires more time to carry out the forecast on the Cray T3E and as in the previous test case, the Cray J90 displays the smallest elapsed time (Figure 5). However, the speedup graph, Figure 6, shows that the Cray T3E is as efficient as the Cray J90, and scales better than theHIVE. If the speedup obtained on the Cray J90 and theHIVE (with one processor as reference) is determined, better scalings on theHIVE compared to the Sesame test problem will be seen. This is due to the improvement of the load balancing associated to the size of the Hurricane Opal problem.

To determine possible causes of the moderate performance of MM5 on the Cray T3E (largest elapsed time), the team estimated how much time it took to execute each section of the code, namely the initialization time (Init. Time), the I/O time and the time spent to carry out all the time steps and radiation calculations (Com. Time); the sum of all these times gives the elapsed time (El. Time). The findings are reported in Table 2, where results show that a large amount of time is spent during the initialization phase on the Cray

T3E. A possible reason for this long initialization procedure may be the large time it takes to retrieve the data. But the team carried out several other experiments and found out that the time to retrieve data was fairly small. In experiments performed in another Cray T3E system (outside GSFC), a different result was obtained [5]. In addition, runs on theHIVE for this test problem show that the initialization step is relatively small. This leads us to conclude that some environment variables on the Cray T3E at NASA GSFC are creating initialization overheads during the execution of MM5. The removal of these overheads will substantially improve the parallel performance of MM5 on the Cray T3E. If the focus is just on the time spent to perform the other operations, results on the Cray T3E and theHIVE are comparable.

**Table 2: Hurricane Opal: overall execution time (in seconds) and times (in seconds) spent on each section of the code.**

		Number of CPUs							
		1	2	4	6	8	10	12	16
Cray J90	El. Time	2800	1434	743	519	400	326	286	235
	Init. Time	63.5	37	23	23	36.5	15	16.5	16
	Com. Time	2728.5	1390	714	488	356.5	303	213.5	214
	I/O Time	8	7	6	8	7	8	6	5
Cray T3E	El. Time			2550	1764	1399	1196	973	774
	Init. Time			706	509	436	404	288	235
	Com. Time			1830	1242	950	778	670	524
	I/O Time			14	13	13	14	15	15
the HIVE	El. Time	6216	3193	1767	1228	995	857	752	640
	Init. Time	163	88	59	50	51	48	50	62
	Comp.	6042	3090	1694	1174	931	794	687	564
	I/O	11	15	14	14	13	15	15	14

Finally, the plots of the Mflops per second and the cost/Mflop as function of the number of processors are presented in Figure 7 and 8. As in the Sesame test case, the Mflops/s is better on the Cray J90 and the Cray T3E and theHIVE have comparable results. theHIVE achieves the best performance versus cost result.




Figure 7. Hurricane Opal: number of floating point operations (Mflops) per second as function of the number of processors.

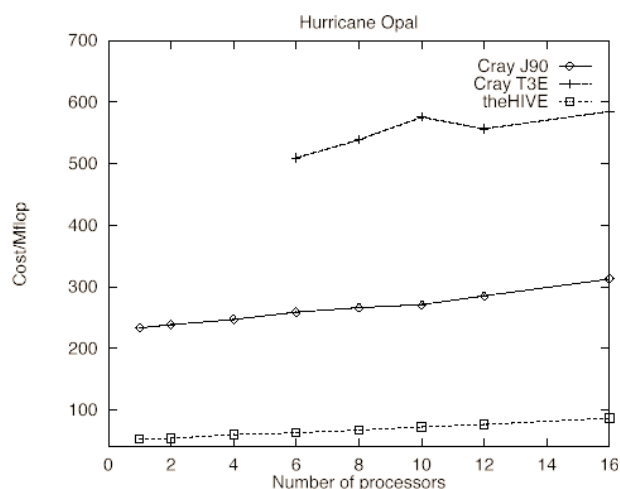


Figure 8. Hurricane Opal: Cost/Mflop as function of the number of processors.

To determine the Mflops/s the team divided the total number of floating point operations by the elapsed time. The number of floating point operations was obtained on the Cray J90 through the `ja` command, on the Cray T3E using the Performance Analysis Tool (PAT) and on theHIVE using the performance monitor `perf` that is a hardware operation counter for Pentium Pro processor.

MM5 was designed to efficiently work on vector and shared memory parallel computers. The efficiency of MM5 on the J90 at NASA/GSFC comes as no surprise. The moderate performance of MM5 on distributed memory (DM) parallel computers (the Cray T3E and theHIVE) does not remove the merit of having a version of MM5 running on such platforms but instead offers the possibility of using large domain problems and a large number of processors. The strategy carried out to obtain a DM parallel implementation of MM5 relies on tools that eliminate the need of rewriting the entire code. The same strategy can be employed on other weather and climate models [5].

## Conclusions

From the performance analysis of the MM5 weather model, results show that on the Cray J90 the scaling is fairly good on both small and large domain problems. MM5 performs well on the Cray T3E and theHIVE when the problem is large. In addition, theHIVE offers the best performance rel-

atively to the cost. On the Cray T3E, it appears that (compared to the other two systems), the initialization time is very high. This has slowed down the overall execution time. The team has not been able to explain initialization overheads since on other Cray T3E systems (outside NASA GSFC), initialization time was far lower.

Learn more about the performance study by emailing any of the authors:

- John Dorband <dorband@gsfc.nasa.gov>
- Jules Kouatchou <kaouatch@pecos1.gsfc.nasa.gov>
- Udaya Ranawake <udaya@neumann.gsfc.nasa.gov>
- John Michalakes <michalak@mmm.mmm.ucar.EDU>

## Acknowledgments and References

The authors wish to express their appreciation to NCAR for providing the MM5 code and are indebted to NCCS for making their exceptional computer environment and resources available. They also thank Milton Halem for his advices and making the project possible.

1. Jimmy Dudhia, Dave Gill, Yong-Run Guo, Kevin Manning and Wei Wang, PSU/NCAR Mesoscale Modeling System Tutorial Class Notes and User's Guide: MM5 Modeling System Version 2, manual available at <http://www.mmm.ucar.edu/mm5/doc.html>.
2. Georg A. Grell, Jimmy Dudhia, David R. Stauffer, "A Description of the Fifth-Generation Penn State/NCAR Mesoscale Model (MM5)," NCAR Technical Note NCAR/TN-398+STR, June 1994.
3. Philip L. Haagenson, Jimmy Dudhia, David R. Stauffer and Georg A. Grell, "The Penn State/NCAR Mesoscale Model (MM5). Source Code Documentation", NCAR Technical Note NCAR/TN-392+STR, March 1994.
4. John Michalakes, "MM90: a scalable parallel implementation of the Penn State/NCAR Mesoscale Model (MM5)," *Parallel Computing*, Vol 23, No. 14, 2173-2186, Dec. 1997.
5. John Michalakes, "Same-Source Parallel Implementation of the PSU/NCAR MM5," *Proceedings of the Second International Workshop on Software Engineering and Code Design in Parallel Meteorological and Oceanographic Applications*, June, 1998, Scottsdale, Arizona, Preprint NASA GSFC/CP-1998-206860, pp.129-139.
6. John Michalakes, "FLIC: A Translator for Same-source Parallel Implementation of Regular Grid Applications," Tech. Rep. ANL/MCS-TM-223, Mathematics and Computer Science Division, Argonne National Laboratory, Argonne, Illinois, March 1997.
7. John Michalakes, "RSL: a parallel runtime system library for regional atmospheric models with nesting," *Proceedings of the IMA workshop Structure Adaptive Mesh Refinement Methods*, Minneapolis, March 12-13, 1997. Also ANL preprint ANL/MCS-P663-0597. ■

# ***Bridging the Gap, an NREN Workshop***

**T**he NASA Research and Education Network (NREN) Project at Ames Research Center hosted a workshop this past August for the High Performance Network Applications Team (HPNAT) and the Networking Research Teams (NRT) to facilitate convergence between the research objectives of the networking technologists and the networking requirements of the applications community. At the workshop, three technology subgroups focused on Quality of Service, Advanced Multicast, and Security in order to prepare the applications developers to take advantage of the new technologies implemented within high-performance networking testbeds (NREN, NISN, ESnet, DREN, vBNS, Abilene, SuperNet, CAIRN).

Case studies were solicited from the applications communities describing the challenges faced in distributing applications across high performance networks and the needs that cannot be met by today's networks. The technology communities working within the NRT and the Internet Security Team were asked to describe what they are trying to provide for applications, and what they think the applications will be able to do tomorrow that they cannot do today.

Representatives from the Joint Engineering Team (JET) and SuperNet communities were invited to assure realism as to what the networking testbeds need to provide in order to support the technology infusions and applications demonstrations. Representatives from the university (e.g., Internet2) community were also invited, to bring in their related activi-

ties and to assure the applicability of workshop results to researchers and applications nationwide.

Participation in the workshop was by invitation only to about 90 invitees selected from the HPNAT/NRT/JET/I2 communities, as well as by selection of invited presentations. The workshop considered the case studies and technology themes in depth, and then broke into subgroups organized to produce specific road maps and collaboration processes by which candidate expected NRT technology outcomes will get infused into JETnet testbeds and targeted HPNAT applications championed by NGI agencies and Internet2 university researchers.

The case studies included:

- Digital Earth
- Digital Video
- Telemedicine
- China Clipper
- the National Partnership for Advanced Computational Infrastructure

Each case study presented an overview, current status (functionality, performance, and limitations), technologies used, and future directions.

A consensus from the workshop - The Next Steps - is forthcoming and will be published in a future issue.

*Excerpted from conference material on the NREN web site. ■*

## ***IAE Launches A New Website***

**J**et Propulsion Laboratory's Interactive Analysis Environment (IAE) Program recently launched a revamped website, now available to the public. The new site describes the Program and allows viewers to keep current with the its latest activities. The intent is to provide broad and rapid access to accurate, easily understood products that convey to the public the significance of mission science results.

IAE, a program built in 1994 upon a group of existing information system research activities, has grown in size and importance over the past five years. This Program applies

emerging technologies to space science missions, with the goal of improving overall productivity.

IAE supports three specific user communities: mission and instrument builders, space scientists using space science data, and the general public. Research activities within the Program provide builders and scientists various capabilities, such as integrating new techniques and interactive distribution.

Learn more about IAE at <<http://www.jpl.nasa.gov/iae>>. ■



# A Glimpse of Supercomputing in the 21st Century

During the week of November 15 Portland, Oregon, hosted SC99, the High Performance Networking and Computing Conference, where the focus was on the impact that high-end technology has had on the rest of the computing world. This impact has produced the trickle-down effect of innovations in a wide spectrum: architecture performance, wide-area networking, visualization, and high volume storage. At SC99, invited speakers, special sessions, and technical papers addressed the issues of trickle-down:

- Is it happening fast enough?
- Do we need to develop more deliberate strategies for technology transfer?
- What societal obligations are associated with technological innovation, and who should assume responsibility for them?

Also explored during the event - how numerical modeling and simulation techniques developed for high-end applications have transformed industrial research and development. In addition, SCinet99 showcased the network technologies of the future and deployed high-speed free-space optics, networking that can now be accomplished literally through the air — without wires or fiber.

By bringing together scientists, engineers, managers, educators, and researchers from all areas of high-performance technology, SC99 provides a powerful arena for exploring the theme of how high-end innovations affect the rest of the computing world.

## Conference speakers

Donna Shirley, former manager of the Mars Exploration Program at the Jet Propulsion Laboratory and the original leader of the team that built the Sojourner Rover, delivered the keynote address on Tuesday, November 16th. She spoke on how creative enterprises and their management must change as high-speed networks and computers become more pervasive. During her address, she presented a system for the management of creative enterprises and showed how it was used to meet the computing challenges of the Mars Pathfinder project and the ongoing Mars Exploration Program. These computing challenges included Pathfinder's application of high performance (for its day) commercial computing technology to a space mission — for the first time. But at the other end of the spectrum the Pathfinder rover made its slow way around Mars using a very low-tech 100 MIP processor. Extrapolating to the future, she described how high performance computing may enable virtual reality to compete with actually going to Mars.

Shirley retired from NASA last year, becoming a speaker and consultant on the management of creative enterprises. In addition to her autobiography, *Managing Martians*, Shirley wrote *Managing Creativity: A Practical Guide to Inventing, Developing, and Producing Innovative Products*.

Invited speakers throughout the five day conference included:

Tom Sterling, of California Institute of Technology and NASA Jet Propulsion Laboratory, who addressed the challenges imposed by the ambitious goal of petaflops computing and discuss the opportunities afforded through alternative device technologies and the HTMT architecture strategy. Detailed findings from the HTMT project will be presented, demonstrating both the feasibility of implementation and practicality of use.

Dona Crawford, director of the Advanced Product Realization Program at Sandia National Laboratories, who provided background for the need to do distance computing, the technology roadmap to achieve the desired capabilities, and the planned methodology to reach 100Gbps by 2004.

Andrew Chien, Science Applications International Corporation Chair Professor at the University of California-San Diego, who described the design and experience of a very large NT cluster, including the core cluster technologies (HPVM) and the system called the NT Supercluster.

Mark Ellisman, Professor of Neurosciences and Bioengineering, Director of the Center for Research on Biological Structure, and Director of the National Center for Microscopy and Imaging Research at the University of California -San Diego, who described the development of novel techniques for 3-D visualization of neuronal structures and modeling of their dynamic properties. He will place special emphasis on examples that involve the application of parallel processing and distributed computing.

Chris Johnson, Professor in the Departments of Computer Science, Mathematics, Physics, and Bioengineering at the University of Utah, who presented new large-scale modeling, simulation, and visualization techniques for applications in computational combustion from the ASCI Center for the Simulation of Accidental Fires and Explosions (C-SAFE) and computational medicine from the Center for Scientific Computing and Imaging. He will also compare and contrast methods in computational combustion and medicine.

John R. Koza, Consulting Professor at Stanford Medical School and President of the Third Millennium Venture Capital Limited who introduced genetic programming; illustrated its application to problems of control, classification, system identification, and design; presented some previously patented analog electrical circuits that have been automati-

cally created using genetic programming; and described the parallel computers (both transputer-based and beowulf-style) that have been used to produce these results.

In addition, State-of-the-Field talks were presented.

Vinton G. Cerf, Senior Vice President of Internet Architecture and Technology for MCI WorldCom, explored how the Internet is likely to evolve over the next four to six years and describe implications for how this evolution might affect scientific computing and other science industries that rely on the Internet.

Salvatore J. Stolfo, Professor of Computer Science at Columbia University, described approaches to substantially increase the amount of data a knowledge discovery system can handle effectively over distributed data sources and speculate about a new generation of KDD/DM systems that might begin to take shape in the next few years.

Greg Papadopoulos, Chief Technology Officer of Sun Microsystems, examined the watershed trend in high performance architecture: supercomputers created from inter-networked high performance, yet commercial, computing platforms. What are the driving forces behind these platforms? How do we expect system architecture, interconnection networks, storage, and software stacks to evolve?

Daniel A. Reed, Professor and Head in the Department of Computer Science at the University of Illinois at Urbana-Champaign, examined some of the reasons why obtaining high performance remains difficult, lessons from the history of high-performance computing, and the challenges for emerging systems. He will highlight approaches that aid understanding and improving performance now, and how things may change during the next five years.

## Webcasts

For the second year, those who could not attend the week-long event could view webcasts of Donna Shirley's keynote address, as well as selected talks: "State of the Art of the Internet," by Vinton Cerf, "Distributed Mining: Problems and Opportunities," by Salvatore Stolfo, Professor of Computer Science at Columbia University, "HPC meets .com: The Convergence of Supercomputing and Super-Internet Architectures," by Greg Papadopoulos, and the Awards session talk, "Augmented Reality and Tele-collaboration with Seas of Cameras and Projectors: Computational Nightmare or Nirvana," by Henry Fuchs of the University of North Carolina.

Additionally, the feasibility of webcasts for the deaf was demonstrated, using sign language and real time captioning as a service on high performance grids. The SC99 webcast team worked with the Trace Center at the University of Wisconsin at Madison on this effort.

## NASA exhibits, papers, and panels

NASA's research exhibit highlighted important high performance networking and computing research projects at various installations. A wide variety of workstation, networked connections to high-end systems, video theaters, virtual real-

ity devices displayed the research and encouraged visitors interaction.

## Ames Research Center

Visitors had the opportunity to grab a computer-generated Earth with their bare hands, perform 'virtual' surgery, or interact with complex molecular structures when they visited the NASA venue in the SC99 Exhibit hall. These and other cutting-edge NASA computer and network technologies were on display during the event. Scientists and engineers from various NASA centers demonstrated and explained their latest computer simulations, ranging from medical and geographical imaging, to advanced human-machine interfaces, aerospace vehicles, galaxy formations and new learning technologies. According to Bill Van Dalsem, deputy program manager of NASA's High Performance Computing and Communications Program, the general public had an opportunity to interact with and understand developing technology and chat one-on-one with researchers.

A variety of collaborative-environment technologies that allow scientists, doctors and engineers to develop new procedures and improve existing ones were on display at the NASA booths. In one demonstration, scientists from ARC showed how highly sophisticated medical imaging combined with high-performance networking can be used to "bring the clinic to the patient." NASA is supporting remote collaborations of doctors at different locations on Earth, preparing to use the technology for spacecraft crews traveling to Mars or other planets where specialists may not be available. This telemedicine technology will allow physicians to consult, diagnose, and plan treatments for patients in real time from a great distance using 3-D images rendered on high-performance computers. This virtual collaborative clinic will help doctors treat astronauts traveling in space and provide care for people in remote locations on Earth. During the demonstration, visitors to the NASA booth will be able to play the role of physician and 'operate' on hearts, skulls and other body parts using this unique software. The NASA Research and Education Network engineered a solution to make distribution and use of the virtual collaborative clinic images possible.

Another demonstration, the virtual mechanosynthesis simulation experiment, will allow users to practice designing models with vibrating and rotating simulated atoms the size of ping pong balls; a form of assembly that is the essence of nanotechnology. Nanotechnology is the control of matter on the nanometer scale, typically from one-tenth of a nanometer to 100 nanometers; a nanometer is one billionth of a meter. Nanotechnology is also the construction and operation of machines on the nanometer scale. Researchers search for "mechanisms capable of placing individual atoms in precisely defined positions. With this technology, researchers can move individual atoms and create plausible atomic designs. A force-feedback arm device even allows users to 'feel' the forces at work between atoms.

## Goddard Space Flight Center

Visitors were able to build their own potato-chip-shaped carbon hydrogen junctions, and even “touch” atoms - the smallest particles of elements — for the first time. An additional feature of the NASA display was the digital Earth immersive workbench developed by scientists and engineers at Goddard Space Flight Center. Workbench users wear stereo glasses to view a 3-D image of the Earth that they manipulate with intuitive hand movements. Digital Earth explorers can hold the Earth in their hands and lift it up to their faces to see natural and cultural forces that affect our planet. Viewers have the option of observing sea surface temperature measurements, weather forecasts or movies of recent earthquakes and other disasters.

## Jet Propulsion Laboratory

In collaboration with the California Institute of Technology, JPL featured the Hybrid Technology Multi-Threaded (HTMT) class architecture for petaflops scale computation. A 3-D walkthrough of the HTMT machine room was demonstrated on several workstations. Showcased hybrid technologies were also on display.

## California Institute of Technology

The Center for Advanced Computing Research at Caltech featured “See the Light,” with live computer demonstrations, technical publications and information on its latest computational research activities.

## Technical papers and panels

NASA was represented in the papers and panels arena by Marjory Johnson from ARC who presented a paper entitled “Using the NREN Testbed to Prototype a High Performance MultiCast Application.” Tom Sterling of Caltech and JPL participated in the panel discussion on large scale computing, “Beyond Grids: Large Scale Computing in a Connected World.”

The supercomputing conference series is an annual event sponsored by IEEE and ACM that began in 1988. The event is organized by a voluntary consortium of persons acting together to advance the science and application of high performance computing and communications technology. SC2000 will be held in Dallas, Texas.

*Excerpted from SC99 conference materials and NASA press releases. SC99 Logo courtesy of the conference web site. ■*



***Through its many outreach programs NASA makes significant contributions to American education systems and to institutions dedicated to improving science literacy, and provides science research results and information to the general public. This newsletter's objectives are to universally educate and provide the community with access to reports of hands-on technology research.***

# History and Radio JOVE Join at NRAO

***James Thieman, National Space Science Data Center, Goddard Space Flight Center***

The Radio JOVE education and outreach project recently emulated radio astronomy pioneer, Karl Jansky, by successfully observing Jupiter using the Radio JOVE receiver kit connected to a reconstruction of the 1932 Jansky telescope at the National Radio Astronomy Observatory (NRAO) in Green Bank, West Virginia. The Radio JOVE project gives schools and the general public the opportunity to build a radio astronomy receiver and antenna and use them to record and analyze the radio emissions from Jupiter and the Sun.

## Presenting JOVE

Project members Chuck Higgins, a Research Associate of the National Research Council, and Len Garica and James Thieman, Goddard Space Flight Center scientists, journeyed to Green Bank together with University of Maryland astronomy student, Albie Davison, and Annapolis High School student, Autumn Thayer, who were working with the project for the summer. There, the group met Richard Flagg, another project member from Hawaii and designer of the Radio JOVE receiver. Together, they gave a presentation about Radio JOVE to the annual meeting of the Society of Amateur Radio Astronomers.



By good fortune the meeting coincided with a time period when Jupiter radio storms had been predicted. The group not only set up the standard Radio JOVE receiver kit and antenna but also connected to the reconstructed Karl Jansky telescope, which NRAO maintains at its facility. The effort was rewarded by the recording of Jupiter radio emissions at NRAO and their verification by simultaneous acquisition by Francisco Reyes, another project member, at the University of Florida Radio observatory.

### Monitoring Galileo

In the Radio JOVE electronic newsletter, Len Garcia encouraged all Radio JOVE teams to try to schedule some

observations in the days before, during, and after the Galileo spacecraft makes its close passes by Jupiter and its moon, Io, during the subsequent months through December 1999. Several Jovian radio storms - that are well placed for Radio JOVE observers - were expected to occur at night while Jupiter is well above the horizon for most US observers. Astronomers from around the world will be observing Jupiter and Io with an assortment of ground-based and space-based telescopes. These observations will be in the radio, infrared, optical and ultraviolet wavelengths.

To participate in the program complete the application form available at the Radio JOVE website. You may also contact J.Thieman via email <thieman@nssdc.gsfc.nasa.gov>. ■



***Through its many outreach programs NASA makes significant contributions to American education systems and to institutions dedicated to improving science literacy, and provides science research results and information to the general public. This newsletter's objectives are to universally educate and provide the community with access to reports of hands-on technology research.***

## Trip to the National Radio Astronomy Observatory

***Autumn Thayer, Annapolis High School, JOVE Participant***



The 140 foot telescope at Green Bank.

**O**ur trip to the National Radio Astronomy Observatory (NRAO), in Green Bank, West Virginia, on July 12-13, was a great success, though we were left quite tired afterwards. The presentation of the Radio JOVE project went very well with the members of the Society of Amateur Radio

Astronomers (SARA). Many were interested and impressed with what we had to offer.

They were especially impressed with the JOVE software. We got many questions in reference to whether or not we sold just the software. Well, we don't, but they were happy to learn that it will soon be available for free downloading from the Radio JOVE website. But above all other interest, SARA agreed with us that it is a good idea to educate and reach out to students about radio astronomy. Hopefully they will be able to help the schools around them become involved with Radio JOVE, and assist them if they ever run into any problems. It is good to have their support since Radio JOVE is very similar to what they do, but with less intricate, and less expensive equipment. Overall, the conference was a success and Dr. Chuck Higgins, Dr. Jim Thieman, and Dick Flagg did such a wonderful job presenting, that we even got some people interested enough to come out and observe with us, at 2 till 6 in the morning. (Talk about interest and dedication!)



Listening for Jupiter at sunrise.

The observations were made on the front lawn of the NRAO beside the replica of the Jansky antenna, which we also used in our observations that cold morning. Besides the Jansky antenna, we also set up three other antennas, which we had brought with us. They were a dual dipole, a single dipole, and a loop antenna made by Jim Gass. At the beginning of our observations, Jupiter was supposed to be low in the sky and the lobe of the dual dipole has a beam pointed directly upward. And so to remedy this problem, we increased the length of cable between the two dipoles, so that we could extend the dipoles one wavelength apart and have the beam pointed more toward the horizon. The funny thing is that we didn't extend the two cables until we had already set up the antenna the night before observing. So it was quite interesting seeing a bunch of men trying to solder cable, in the cold, in the middle of a field. We actually had to use a lighter to help keep the connection warm enough to solder, because it was so cold.



Everybody soldering cable.

Well, luckily, the extension of cable worked well and we received a lot of radio emission from Jupiter. We also got great results for all the other antennas, including the Jansky antenna. We picked up some L-bursts, and actually quite a few S-bursts. Another neat thing about our data is that at the same time we were receiving lots of data, the people down at the University of Florida were receiving data too. So it was confirmed; it was definitely Jupiter we were hearing.

So this trip had many successes. We got some of the SARA members interested in the project, and we got a lot of support from them. Hopefully this will continue in the future. We also got great data, but more importantly, we proved that Radio JOVE does work. It actually works very well, and we have some of the data to show for it. But with life, and especially in science, you live and learn. We have learned that you must always check the recording level of the recorder while you are observing, or you might be left with just background noise, like we were. So this is a warning to any who might make the same mistake: Be sure to check that the recording level is not too high. Fortunately, one of our recorders was set correctly and so we retained some of the data. But, all in all, the trip was very successful, and for those that were there it was very exciting to hear Jupiter through all our Radio JOVE antennas. It was an excellent trip, well worth suffering the cold at 2 o'clock in the morning.



Goddard group beside the Reber telescope.

To participate in the program complete the application form available at the Radio JOVE website. You may also contact J.Thieman via email [<thieman@nssdc.gsfc.nasa.gov>](mailto:thieman@nssdc.gsfc.nasa.gov). ■



***Through its many outreach programs NASA makes significant contributions to American education systems and to institutions dedicated to improving science literacy, and provides science research results and information to the general public. This newsletter's objectives are to universally educate and provide the community with access to reports of hands-on technology research.***

## ***NASA Sponsors “Mars Millenium” Internet Webcast***

**D**esigning a Mars community for human beings was featured during a NASA-sponsored educational one-hour “webcast” on September 28 through Quest at Ames Research Center (ARC). Four NASA scientists participated in a round-table discussion during the webcast to answer students’ questions submitted during a simultaneous “web chat,” as well as by students at the event site. During Internet chats, students used computers to converse with mentors by typing questions and reading responses and dialogue via the World Wide Web. The webcasts included a live, moving video picture and sound that were sent to computers via the Internet.

According to Sandy Dueck, webcast moderator for ARC, the scientists discussed four questions about Mars during the webcast:

- What is Mars like right now?
- Why are we so interested in Mars?
- What are the current missions planned for Mars?
- When and how will human beings explore Mars, and what might they do there?”

The webcast was incorporated into the Mars Millennium educational program that is endorsed by the White House. Dueck explained. “This is an Internet educational activity that educators across the country will use,” she said.

The webcast was hosted from ARC’s Main Auditorium. There, Ames researchers Christopher McKay, a planetary scientist whose current research focuses on the evolution of the Solar System and the origin of life; and his colleague, Kelly Snook, who analyzes Mars data and is working on a possible mission to Jupiter’s moon, Europa, made their presentations. Also participating were Robert Anderson, who is the science mission planner for the Mars 2001 mission and the Mars Pathfinder outreach coordinator, and David Seidel, an outreach supervisor and pre-college programs officer. Both are from Jet Propulsion Laboratory.

Four classrooms of students from three San Francisco Bay Area schools, grades 2 through 8, also participated on site at ARC in the round-table discussion with the scientists. The schools were:

- Monarch Montessori, Sunnyvale, California
- Hayward Project School, Hayward, California
- Toyon Elementary School, San Jose, California

The Mars Millennium Webcast project is one of many Internet offerings from NASA’s Quest Project. These on-line, interactive projects connect students with NASA employees and are designed to inspire young people to pursue careers in high technology. ■



***NASA's wealth of technology is being re-used in the fields of medicine, industry, education, and by the military to develop products and processes that benefit many sectors of our society. Spinoff applications from NASA's research and development programs are our dividends on the national investment in aerospace.***

## ***NASA Technology Used to Track Threatened Birds***

**Y**ou can watch birds fly right on your computer monitor! You can also learn how scientists study the flight patterns and migration routes of threatened or endangered birds, including harpy eagles in Panama and Venezuela, African eagles, cranes in Siberia and Canada, steppes and golden eagles in Mongolia, and fish eagles in Madagascar.

In 1990 the Patuxent Wildlife Research Center joined forces with the Goddard Space Flight Center (GSFC) in a project that uses telemetry to discover the winter home of a dwindling population of Siberian cranes. Since that initial project, additional projects have emerged to study other endangered birds. Subsequently, a web site was created at GSFC to allow the public to follow the tracking activities. On the site, the program is described in detail at three links: the satellite direct readout, birdtracks, and birdworld.

The satellite direct readout provides an overview of the program, the agencies involved, and how a bird is actually tracked. This link explains how the Patuxent scientists study the birds' habitats and prepare to track the birds by fitting each with tiny transmitters. It also describes how NASA processes and correlates the data into bird locations.

In addition to the Patuxent Wildlife Research Center and NASA, Service Argos plays a key role in the feasibility of the

program by enabling access to the satellite data from transmitters around the world. At Service Argos, the data is collected, processed, and passed electronically to scientists at NASA.

The birdtracks link provides a graphical demonstration of the birds flying along their flight path, animated through time. In order to present this demonstration, NASA programmers must electronically pinpoint bird locations on digital maps, using their experience in geo-registration of satellite data. NASA has developed software to automatically locate the bird transmitters and represent each location on the maps.

The birdworld link features a slideshow of the birds tracked in the project, as well as their natural habitats. You can also read actual trip reports and published articles that describe the unique studies being developed to protect these endangered birds.

*Excerpted from material provided by Margaret Williams, NASA web developer and website content author. Many thanks to Dr. David Ellis, Patuxent research scientist and principal investigator.*

Learn more about endangered birds at  
<[http://sdcg.gsfc.nasa.gov/ISTO/satellite\\_tracking](http://sdcg.gsfc.nasa.gov/ISTO/satellite_tracking)>. ■